

EXHIBIT 11

_docx

Main document changes and comments

Page 10: Commented [1]	Joelle Pineau	2/15/2023 5:07:00 PM
-------------------------------	----------------------	-----------------------------

Lack of coordination with Meta product strategy may be a miss (for us).

Page 10: Commented [2]	Joelle Pineau	2/15/2023 5:08:00 PM
-------------------------------	----------------------	-----------------------------

Not a blocker, but mentioning for completeness.

Page 11: Commented [3]	Natalie Hereth	2/15/2023 10:02:00 PM
-------------------------------	-----------------------	------------------------------

similar to Joelle's comment above, do we want a question on what the impact of this release could have on future product plans?

Page 11: Commented [4]	Faisal Azhar	2/16/2023 9:17:00 AM
-------------------------------	---------------------	-----------------------------

We had that Q in the key decision section and later removed it to limit the discussion to Key decisions 1A and 1B. I think that impact on future prod plans will organically come up.

Page 12: Commented [5]	Natalie Hereth	2/15/2023 10:01:00 PM
-------------------------------	-----------------------	------------------------------

is this still accurate?

Header and footer changes

Text Box changes

Header and footer text box changes

Footnote changes

Endnote changes

[DECIDE] Go To Market Review (GTM):
Genesis (TBC Name) Gated Open-source Release

Document Purpose: We want to announce through Meta AI social channels, blog, and limited press outreach a gated release of the model FAIR LLM, targeting **22-Feb**, with a bespoke license for research usage only¹. The plan is to release; (1) a research paper on ArXiv to show SOTA results against similar size models, (2) gated open-source model weights of sizes 1.5B / 13B / 33B / 65B for research exploration reducing carbon footprint by enabling researchers to re-use the models, (3) GitHub repository to allow researchers to load model and perform inference, (4) model cards for understanding, sharing and improving machine learning models, (5) a google form to decide when and who gets access to our models. In this GTM review, we want to decide to Go/NoGo with our external release plan with the risks and mitigations in place due to the increased scrutiny and attention on large language models (LLMs).

1. Key Decisions for GTM Meeting

- A. Does the quality of research results and the rigor of the release plan meet our high bar to do an external release?
 - a. If not, then why?
 - b. If yes, then
- B. Are we comfortable with the level of risk and the mitigations that have taken place to launch FAIR LLM externally?
 - a. If not, then why?
 - b. If yes, then

2. Situational Context

With the increased research advancement in LLMs, we want to continue to position FAIR as a leader in this space. We believe that by announcing FAIR LLM, there is an opportunity to showcase our innovation in the impact smaller models bring to the LLM space is valuable in terms of lower cost deployment in products, re-training and fine tuning towards specific downstream applications and reducing the carbon footprint to re-use the models for research. By releasing the smaller models, we democratize access and study of models for academics who want to do research on LLMs and do not have the computing infrastructure to train their own. Research on smaller LLMs is a really hot topic (see HF CTO tweet) and we are expecting a lot of competition. We are currently well positioned to show that FAIR is now competitive in this space.

We will be the first to inform the community that our smaller size models are competitive with the SOTA large language models, thereby avoiding being scooped by DeepMind and Stability AI. Our research work is building upon the recent research from DeepMind (GPT, PaLM and GPTNeo). Hence, we have strong reasons to believe other external groups (DeepMind, Stability AI - will open-source their model StableGPT¹, very soon) are also working on the same research work and thus we want to be the first to inform the community that smaller sized models (<=70B) can be competitive with existing larger models (6B - 540B).

FAIR LLM Competitive Advantages related to competition (GPT-3, PaLM, Chinchilla, Gopher):

We present a brief summary of our model quality across four axes in the table below and refer to our paper (WIP) for details.

¹ Stability ai will release StableGPT and they are also done with their training, so they could do a release any moment which puts us at high risk of being scooped by the competition.

Action Items:

- **[Action Item]:** Risk number 3- possible bias / representation on the dataset. Would it be possible to add to the Appendix the evaluations of the model and the dataset? *What we tested against vs. not tested.*
Done. Added table 4 in section.

- **[Action item]:** Make it clear on what's the criteria of who is going to get access to the model and share back with the community who got access to it.
Done, criteria added in section.

- **[Action item]:** Request access to the blogpost draft. Ashley will share it.
- **[Action item to Jon]:** to brief MZ's team on the GTM
- **[Action item]:** what do we expect the team to react from their personal accounts on the launch? Social moderation and social guidelines and communication to the teams. Alman Farooq

In progress, doc.

No	Axes	Summary
1	Performance on benchmarks	<ul style="list-style-type: none">See table below
2	Reproducibility	<ul style="list-style-type: none">Our models are trained on only publicly available data so making our work reproducible and compatible with open-sourcing values.
3	Deployability	<ul style="list-style-type: none">Our models are quite smaller than LLMs so they are significantly easier to deploy at scale on even a single-GPU
4	Bias, Toxicity and Misinformation	<ul style="list-style-type: none">Our model bias, toxicity and misinformation is evaluated on standard benchmarks and is similar to GPT-3 and OPT LLM.

Table 1 - Summary of Genesis Model Competitive Advantage

We present our results on eight standard common sense reasoning benchmarks in the below table. Our 65B model outperforms Chinchilla-70B on all reported benchmarks and it surpasses PaLM-540B everywhere but on BoolQ and WinoGrande.

Models	Number of Parameters	Reasoning tasks								
		BoolQ	PIQA	SIQA	HellaSwag	WinoGrande	ARC-e	ARC-c	OBQA	COPA
GPT-3	175B	60.5	81.0	-	78.9	70.2	68.8	51.4	57.6	91
Gopher	280B	79.3	81.8	50.6	79.2	70.1	-	-	-	-
Chinchilla	70B	83.7	81.8	51.3	80.8	74.9	-	-	-	-
PaLM	62B	84.8	80.5	-	79.7	77	75.2	52.5	50.4	93
Cont-PaLM	62B	83.9	81.4	-	80.6	77	-	-	-	-
PaLM	540B	88	82.3	-	83.4	81.1	76.6	53	53.4	93
Genesis	7B	76.5	79.8	48.9	76.1	70.1	76.7	47.6	57.2	93
	13B	78.1	80.1	50.4	79.2	73	78.1	52.7	56.4	94
	33B	83.1	82.3	50.4	82.8	76	81.4	57.8	58.6	92
	65B	85.3	82.8	52.3	84.2	77	81.5	56	60.2	94

Table 2 - Summary of Genesis Model Performance on Reasoning tasks

We present our results on Massive Multitask Language Understanding (MMLU) in the table below. We observe that our 65B model is behind both Chinchilla-70B and PaLM-540B by a few percent in average, and across most domains. A potential explanation is that we have used a limited amount of books and academic papers in our pre-training data, i.e., ArXiv, Gutenberg and Books3, that sums up to only 177GB, while these models were trained on up to 21B of books. This large quantity of books used by Gopher, Chinchilla and PaLM may also explain why Gopher is outperforming GPT-3 on this benchmark, even though it is comparable on most of the other benchmarks.

Models	Number Of Parameters	MMLU Benchmark			
		Humanities	STEM	Social Sciences	Other
GPT-3	175B	40.8	36.7	50.4	48.8
Gopher	280B	56.2	47.4	71.9	66.1
Chinchilla	70B	83.6	54.9	79.3	73.9
PaLM	62B	59.5	41.9	62.7	55.8
PaLM	540B	77.0	55.6	81	69.6
Genesis	7B	34	30.8	38.3	38.1
	13B	45	36.2	53.8	53.6
	33B	55.8	46.2	66.7	63.9
	65B	61.8	51.9	72.9	67.7

Table 3 - Summary of Genesis Model Performance on MMLU benchmark

We present our results on bias against GPT-3 and OPT in the table below. Note that lower value is better indicating lower bias and bold show winner against each category. Overall our model FAIR LLM is slightly better compared to GPT-3 and OPT.

No	Category	FAIR LLM	GPT-3	OPT
1	Gender	68.3	62.6	65.7
2	Religion	78.1	73.3	68.6
3	Race/Color	58.7	64.7	68.6
4	Sexual orientation	79.8	76.2	78.6
5	Age	70.1	64.4	67.8
6	Nationality	59.7	61.6	62.9
7	Disability	71.7	76.7	76.7
8	Physical appearance	71.4	74.6	76.2
9	Socioeconomic status	72.1	73.8	76.2
FAIR LLM Average		66.2	67.2	69.5

Table 4 - Summary bias of our model output against GPT-3 and OPT

3. External Release Artifacts

As part of the FAIR LLM release, we enumerate in the below table the release artifacts versus owner, description, approval processes, timeline and status.

No	Artifact	Owner	Description	Approval processes	Timeline	Status
1	Paper	Faisal Azhar	<ul style="list-style-type: none"> Research paper (WIP) external release. 	<ul style="list-style-type: none"> Publication review (~7 business days for non-sensitive) 	~17-Feb	In progress
2	Model & weights	Faisal Azhar	<ul style="list-style-type: none"> Pre-trained model with weights to be released, hosted on AWS S3. 	<ul style="list-style-type: none"> SRT covers legal, policy & IP review including dataset 	~17-Feb	Approved
3	GitHub repo (evaluation)		<ul style="list-style-type: none"> Evaluation code repo containing methods to load the model checkpoint and perform inference. We will NOT release training code externally, that will only be available for internal use - but we will share this internally. 	<ul style="list-style-type: none"> Open sourcing project workflow Trademark review targeting to have a response from TR by 02/20. 	~17-Feb	Redacted - Privilege
4	Model cards	Aurelian Rodriguez	<ul style="list-style-type: none"> Documentation for AI models 	<ul style="list-style-type: none"> Review with XFN. 	~17-Feb	In progress
5	Meta AI blog, social amplification, and media engagement	Ashley Gabriel Alyssa Newcomb Oliver Libaw Faisal Azhar Aiman Farooq	<ul style="list-style-type: none"> Comms narrative 	<ul style="list-style-type: none"> Comms plan approval 	~17-Feb	In progress
6	FAIR LLM request form	Faisal Azhar	<ul style="list-style-type: none"> A request form (to be updated) for researchers to express interest to access our models. 	<ul style="list-style-type: none"> Privacy LAMA (~4 wks) covers user data collection for intake form. 	~17-Feb	Redacted - Privilege

4. FAIR LLM External Release Approach

We are targeting 22-Feb to announce FAIR LLM externally. We will use a brief Meta AI blog, light social amplification on Meta AI channels, and select AI-focused media outreach to land a narrative on the importance of open research narrative in the fast moving field of AI. In our reactive messaging, we will make it clear this is not Meta's response to ChatGPT and Bard. Alongside, we'll publish our research paper on ArXiv which will focus on showing that smaller and moderate size models can be further improved over SOTA in LLMs (Churchilla by DeepMind 2022). The paper will include a link to a google form which will ask basic user information, intended use and acceptance of a bespoke license to get access to our models (form). Each access request to our model will be reviewed by our research lead (Armand Joulin) before granting based on intended use (e.g., research, commercial, other) the actual access to our models.

Criteria:

- We will release our models under a bespoke license to focus on certain research use cases, and we will also take additional precautions by gradually rolling out initial access to our models via a prioritization system. Requesters who:
- provide links to prior publications;
 - have an email address matching an academic institution or industry research lab; and
 - provide extensive details on the intended research use cases


are highly prioritized for initial access to our models. Each request will be manually screened to verify the consistency and accuracy of information provided in the request form before the requestor receives a link to download the model.

[a] Research Objective:	To advance research in LLMs by creating a series of language models that achieve the best possible performance at each size by training on more tokens than what is typically used, and democratize access and study of models by gated release to the research community.
[b] Key Messages:	<ul style="list-style-type: none">• As part of Meta's commitment to open science and transparency, we are releasing Genesis in order to democratize access to large language models and help advance the work in this subfield of AI.• With the FAIR LLM, researchers have found that smaller models, trained on more tokens, can compete well with larger models on existing benchmarks.
[c] Primary audience:	Highly focused on AI practitioners and AI Academics (detailed profiling here). As a secondary audience, we'll indirectly reach key influencers of public opinion and policy since due to increased attention with LLMs.
[d] Success metrics:	Research: Have major researchers of the space comment on the model and generate an important number of citations (100+ within a year), see an important (3000+ within H1) demand for accessing to the gated model. Marketing/Press: 70% neutral sentiment across channels (social, press), with inclusion of our open research/transparency message.

5. Risks & Mitigations

Risk Details		FAIR LLM Mitigations		Severity (Low, medium, high/ blocker)	
Redacted - Privilege		Mitigation feedback	FAIR LLM Research Effort & Status	Mitigation Owner & XFN Risk Issuer	ETA
		Aligned	Completed	Owner: Research XFN: Legal & Policy	14-Feb

Redacted - Privilege			
Aligned	Completed	Owner: Research XFN: Legal & Policy	14-Feb



Redacted - Privilege

HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY

<p>Policy</p> <p>A gated release for Genesis supplements any licensing terms by reducing the probability of misuse/abuse of the model by any downstream bad actors for</p> <ul style="list-style-type: none">• unauthorized public demos hosted elsewhere;• scams or other bad activity;• impersonation of a natural person online (i.e., online "bots");• retraining on hate speech, including for online harassment or bullying; and• generation of misinformation <p>There are also Policy/Comms benefits to gating Genesis, which will generate:</p> <ul style="list-style-type: none">• valuable metrics about the quantity and sources of demand for the model;• evidence for how our model is being used to advance science and the social good; and• consistency in our messaging about maintaining leadership on AI by sharing our research responsibly.	Model gating to legitimate researchers.	Aligned	Completed	Owner: Research XFN: Policy	14-Feb
	<p>Media and public perception that the FAIR LLM is Meta's response to ChatGPT and BARD.</p> <p>We will take a humble approach with the announcement of the Genesis models, leveraging language that is explicitly technically focused, making it clear that this is not a consumer message and is positioned as a research advancement for SMEs.</p>	Aligned	Completed	Owner: Comms XFN: Research, legal policy	17-Feb

RAI assessment on FAIR LLM Gated Release and Technical Artifacts

Most of the RAI associated risks come from the possible interpretation of the public when comparing this model with models that are more product oriented (like chatGPT). There are also competing risks related to RAI expectations for research (such as being open and transparent) and RAI risks of releasing a model that could potentially be misused.

Below is the list of potential risks and proposed mitigation strategies:

No	Risks	Level	Possible mitigations	Blocker
1	Open Science: Since the model is presented quite transparently, the risk of being accused of not being open is relatively low. While a fully openly available model would be	Low	Communicate honestly about what is open, focus on reproducibility and access to a broader base of practitioners due to the portability of the model	No, implementing mitigations

	preferable to the open science movement, Genesis is releasing the source code, weights, and gated access, so overall the risk of not being open is small.			
2	Environmental impact: while ML models are known to use a lot of energy resources, causing CO2 emissions, this model focuses on the efficiency of training, so it's a good example of how effective methods can lower significantly the environmental impact	Low (may be a positive impact)	Make sure to make environmental impact analysis and publish a comparison with more resource hungry models. Avoid hyping claims since we don't know how it's going to be used later	No implementing mitigations
3	Possible bias or lack of representation in the original training dataset	Medium	For concerns of possible bias in the training data, one could argue that bias has been measured at the model output level, showing lower or similar level of bias as other models. In particular, the PaLM paper has a good dataset analysis (Appendix C). For concerns of representation, I think this is a reasonable concern. It is true that the paper doesn't publish nor cites representation distribution in the original data. It would be nice to have this evaluation done across multiple axes, not just to highlight the lack of representation but also to better understand what is being represented and how. To mitigate this risk, a possible argument is that this type of analysis is a research project in itself and more work is planned	No implementing mitigations
4	Transparency & Reproducibility	Low/Medium	While this project is releasing the source code, opening access and using public data, it is always tricky to publish comprehensive evaluations and explain mitigations convincingly enough. In this case, there is a risk of being accused of cherry picking metrics that make the model look better than others. To mitigate this, I would suggest to add a paragraph at the beginning of the RAI section, including an explanation of why and how specific evals were chosen over other available methods.	No implementing mitigations
5	RAI expectations related to product launch	Medium-High	It is likely that, when the WA demo/prototype is launched, the public will associate this model with the performances of the chatbot. While there is going to be communication and mitigations applied directly to the chatbot, it would be important to establish the distance between this research and a product prototype (such a chatbot). To do that add a very complete and interpretable model card for the base model (this one), including recommendations or highlights for a conversational type of training or other use cases.	No implementing mitigations

6. Pros & Cons of Launching Now

Pros	Cons
------	------

<p>1. We will be the first to inform the community that our smaller size models are competitive with the SOTA large language models, thereby avoiding being scooped by DeepMind and Stability AI (see Yann LeCun's recommendation in SRT).</p> <p>2. Democratizing access and study of models, our models are quite smaller than LLMs so they are significantly easier to deploy at scale on even a single-GPU. This is a really hot topic (see HF CTO Wesley) and we are expecting a lot of competition. We are currently well positioned to show that FAIR is now competitive in this space.</p>	<p>1. High likelihood of this release being perceived as Meta's response to chatGPT and BARD leading to inevitable comparisons similar to our own FAIR internal researcher comments.</p> <p>2. If there is a negative press cycle, it might negatively impact other planned announcements regarding our responsible AI. For example, COv2 dataset announcement is slated for 28-Feb. Any other?</p>
--	---

Commented [1]: Lack of coordination with Meta product strategy may be a miss (for us).

Commented [2]: Not a blocker, but mentioning for completeness.

(optional read)

Appendix & Footnotes

Content Details:

- 1. Who are we (Repatee & XFN Team)
- 2. Important Additional Documents
- 3. Research Detailed Roadmap
- 4. SRTs & LAMAs Tracker
- 5. FAQ
- 6. Footnotes

Who We Are (Repatee/XFN Team)

- FAIR LLM:
 - Armand Joulin - Director Research Scientist
 - Aurelien Rodriguez - SW Eng Managers
 - Guillaume Lample - Research scientist
 - Timothee Lacroix - SW Eng
 - Thibaut Lavril - Research Eng
 - Gautier Izacard - Research assistant
 - Hugo Touvron - Research scientist
 - Edouard Grave - Research scientist
 - Xavier Martinet - Research Engineer
 - Marie-Anne Lachaux - Research Engineer
 - Baptiste Rozière - Research Scientist
- XFN:
 - Ashley Gabriel - Comms lead
 - Ana Paula Kirschner Morfarij - Marketing lead
 - Harrison Rudolph - Policy lead
 - Ahuva Goldstand - Legal lead
 - Mallika Mathotra - Privacy lead
 - Cristian Cantlon Ferrer - RAI lead

- o [Eleonora Pressani](#) - RAI lead
- o [Susan Epstein](#) - Civil Rights lead

Important Additional Documents

- [FAIR LLM gated open-source release plan](#)
- [RAI Question/Comments](#)

Research Detailed Roadmap

For context, FAIR needs a set of core models which can be easily leveraged to efficiently develop future research work. Hence, we developed FAIR LLM which is a general purpose large language model trained by FAIR researchers with a goal to provide internally, a simple and hackable codebase for research - improving usability to make the models easy to query, fine-tune and integrate in existing codebases (see [Roadmap](#)). We want to open-source the FAIR LLM model (1.5B, 13B, 33B, and 65B), weights and Github code repo for research usage only to highlight the capabilities of our model on several downstream Natural Language Processing (NLP) tasks and publish a paper on ArXiv. The FAIR LLM model is trained on a subset of the pre-approved [Zucchini dataset](#) and the technical difference between FAIR LLM vs. Zucchini is [here](#).

Date	Deliverables
03-Nov	<ul style="list-style-type: none">Complete training of 1.3B model on 300B and 1.4T tokens
15-Nov	<ul style="list-style-type: none">Complete training of 13B model on 300B and 1.4T tokens
05-Dec	<ul style="list-style-type: none">Complete training of 30B model on 300B and 1.4T tokens
17-Feb	<ul style="list-style-type: none">Complete training of 65B model on 300B and 1.4T tokens
17-Feb	<ul style="list-style-type: none">Complete paper and all other release artifact updates.

List of SRTs/LAMAs in Place

Review Purpose	Date Filed	Status	Date Completed
Publication review	10-Feb	In Progress	
SRT	10-Jan	Approved	14-Feb
Privacy LAMA	25-Jan	In Progress	

FAQ

1. Are we doing a gated model release including paper and all model sizes to be open-source under a bespoke license for research usage only?

Yes, Gated code & weights [similar criteria to OPT-175B]. Paper open for all (e.g. arXiv). License: primary is for research usage only.

2. Are there any plans to make the model commercially available?

See point #1. Beyond this release, we will deploy a different version of this LLM model in our own production and that different version will not be OSS. Also, the product version of LLM model will undergo lots of fine-tuning, eng work, integrity, etc., hence will be very different and more powerful than the model in this release.

3. Are we doing a demo with the same model?

No demo at this time. As per point#2, a different version of this model (after improvements) will be integrated into WA for deployment there, with its own production launch review process.

Commented [3]: similar to Joelle's comment above, do we want a question on what the impact of this release could have on future product plans?

Commented [4]: We had that Q in the key decision section and later removed it to limit the discussion to Key decisions 1A and 1B. I think that impact on future prod plans will organically come up.

4. Do we have alignment at exec level as if we put out our model/paper now in the LLM space it will likely be seen as our response to ChatGPT and BARD?

Meta executive leadership has a heads-up and we should not call this an alignment or approval. We need to put in front of them our specific plan and then get their approval. We need to explain risks & benefits as part of that conversation.

Commented [5]: is this still accurate?

Footnotes